View metadata, citation and similar papers at core.ac.uk

brought to you by CORE

provided by CiteSeerX

PRIMA FACIE OBLIGATION AND DOING THE BEST ONE CAN

By: MICHAEL J. ZIMMERMAN

Zimmerman, Michael J. "Prima Facie Obligation and Doing the Best One Can," *Philosophical Studies*, 78 (1995): 87-123

Made available courtesy of Springer Verlag.

The original publication is available at www.springerlink.com

*****Note: Figures may be missing from this format of the document**

Article:

I believe that one ought, morally, to do the best one can. This is, of course, controversial. I don't propose to debate the issue here, however, but wish rather to address a related one: how such a theory of absolute obligation — a "maximizing" theory — might accommodate prima facie obligation. My sense is that many believe that maximizing theories cannot accommodate prima facie obligation, and that this has led some maximizers to reject the concept of prima facie obligation while leading some of those who appeal to this concept to reject maximization. I shall suggest, on the contrary, that maximizing theories are well-equipped to accommodate prima facie obligation; I shall do this by proposing an analysis of absolute obligation in terms of certain concepts and then using some of these concepts to analyze prima facie obligation. My purpose is threefold: to enhance our understanding of prima facie obligation; to show that appeal to prima facie obligation is not a good reason for rejecting maximization; and to rehabilitate the concept of prima facie obligation in the eyes of maximizers.

1. ABSOLUTE OBL1GATION

The main claim of the maximizing theories I have in mind may be put as follows:

(I) person S ought absolutely [as opposed to prima facie], at time T in world W, to do act A at time T' iff[1]
- (1) there is a world W' such that W' is accessible to S from W at T and S does A at T' in W';
- (2) there is a world W" such that W" is accessible to S from W at T and S does not do A to T' in W"; and
- (3) for all worlds W" such that W" is accessible to S from W at T and S does not do A at T' in W", there is a world W' such that:
  - (a) W' is accessible to S from W at T,
  - (b) S does A at T' in W', and
  - (c) the value for S at T of W' is greater than the value for S at T of W".[2]

Here (and below) T is to be understood as being no later than T'. Also, as will become evident later, A is to be understood as being the sort of thing that can be done more than once and by more than one person.[3]

What (I) says is roughly this: S ought to do A if S can do A, S can refrain from doing A, and any way of refraining from doing A is worse than some way of doing A. (This is rough: "way" is here intended to cover not just the manner but the circumstances in which A is either done or not done). Such an account of obligation is naturally allied with the following sort of diagram:

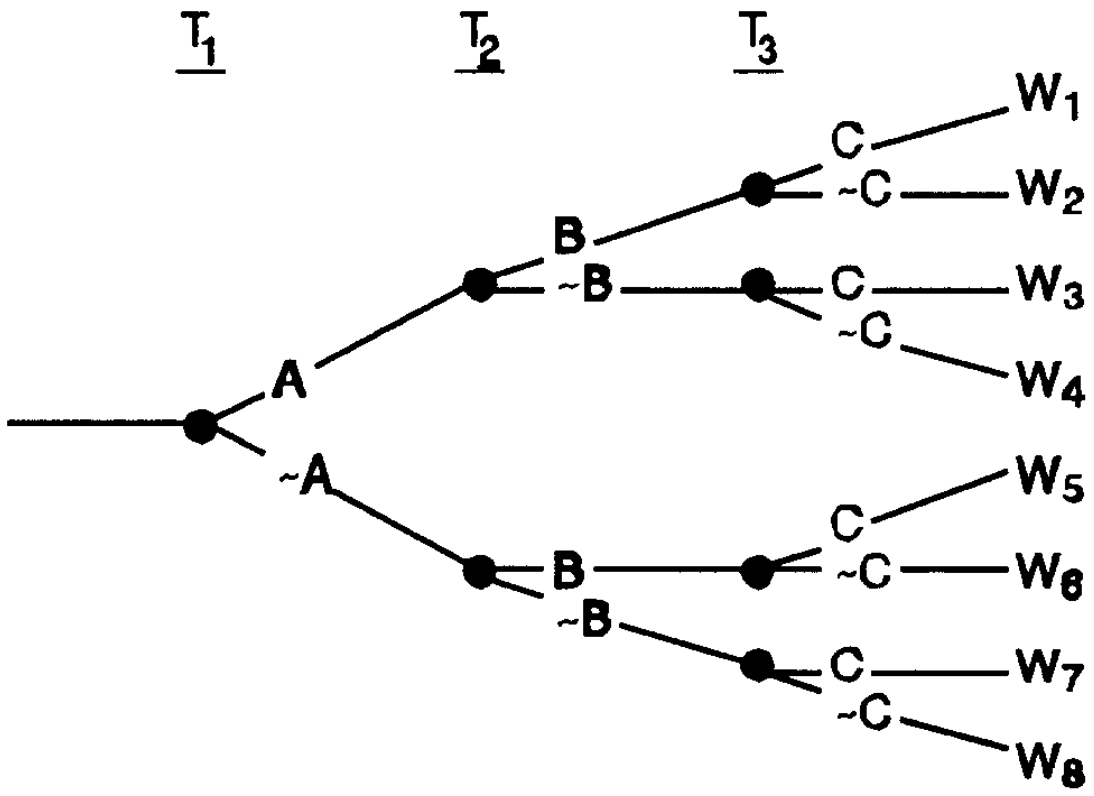

This diagram depicts the following case: at T1 the agent, S, can do each of the following: A at T1, at T1 (that is, S can at T1 refrain from doing A at T1), B at T2, NB at T2, C at T3, and NC at T3. It also depicts the fact that at T1 S can perform any combination of these (except for the mutually incompatible combinations of A and ,A, B and NB, and C and ,C). Of course, this diagram, like all such diagrams, has severe limitations, and the case it depicts is therefore vastly oversimplified. In any actual case, the agent will have many immediate options, not just those of performing a certain action and refraining from performing it; in particular, there will be a variety of ways of refraining from performing it. Still, the diagram suffices as it stands to illustrate what S's obligations are, according to the account given in (I). Let us suppose that the value of W1 (for S at T — more on this in a moment) is greater than that of each of the other worlds. Then the diagram tells us, among other things, the following: S ought at T1 to do A at T1; S ought at T1 to do B at T2; and S ought at T1 to do C at T3. Each of these is true because each of A and B and C is featured in W1. (It is true, of course, that each of A and B and C is also featured in some world inferior to W1, but that is immaterial to their being obligatory). The account given in (I) could be supplemented in obvious ways to furnish us also with the following claims, when applied to the diagram: S ought at T1 not to refrain from A at T1; S ought at T1 both to do A at T1 and to do B at T2; S ought at T1 neither to refrain from B at T2 nor to refrain from C at T3; and so on. I shall not provide such supplementation here, however.

It is natural to associate maximizing theories with consequentialism,[4] but it is not necessary to do so. Although the relevant value that distinguishes worlds may indeed be intrinsic value (as consequentialists claim), there is nothing in (I) that requires this. It is for this reason that the values of W' and W" in (I) are said to be the values that they have for Sat T. If the relevant value is agent- and time-neutral, as intrinsic value is, then such relativization is insignificant and can be ignored, inasmuch as the values for S at T of W' and W" are the values that these worlds have for all possible agents at all possible times. But (I) allows for the possibility that the relevant value is not intrinsic value and also that the relevant value is significantly agent- or time-relative. For example, Ross's sort of pluralistic theory, according to which what is absolutely obligatory is a function of several independent values, can be plausibly interpreted as a maximizing theory; in this case, the relevant value of a world will itself be composed of these various independent values, at least some of which are significantly agent-relative, as we shall see. In fact, there would seem to be nothing in (I) that is inconsistent with Kant's

theory, which is in some ways the paradigm of nonconsequentialism. Suppose that, in the diagram, each of A, B, and C would constitute compliance with the Categorical Imperative while none of ,A, ,B, and ~C would. A "minimal" Kantianism might declare any failure to comply with the Categorical Imperative categorically wrong and thus deem W1 infinitely superior to any of the other worlds, so that (as we might put it) it has a value of 1 while all the others have a value of 0.[5] A somewhat richer Kantianism might declare every noncompliance as bad as any other but allow for them nonetheless to "add up", in which case the ranking would be this: first, W1 (no noncompliances); second, W2, W3, and W5 (one noncompliance each); third, W4, W6, and W7 (two noncompliances each); and last, W8 (three noncompliances). A still richer Kantianism would declare some noncompliances worse than others, and this would allow for a still finer ranking of the worlds. This way of picturing matters is also compatible with saying (as I think almost all Kantians would) that it can sometimes happen that a person can avoid someone else's noncompliance only by way of his own noncompliance. (Think of Bernard Williams's celebrated case of Jim and the Indians; Jim cannot get the captain not to kill all the Indians except by killing one himself.[6] A Kantian would say: either Jim will fail to comply with the Categorical Imperative or the captain will). In such a case, a Kantian would say that it is obligatory for one to comply, even at the cost of the other's noncompliance. This requires that the valuation of worlds be agent-relative, in that S's noncompliance has a special disvalue relative to S that it does not have relative to someone else.

The account of unconditional absolute obligation given in (I) is naturally accompanied by the following account of conditional absolute obligation:

(II) S ought absolutely, at T in W, to do A at T', on the condition that p is true, iff
- (1) there is a world W' such that W' is accessible to S from W at T and S does A at T' in W' and p is true in W';
- (2) there is a world W" such that W" is accessible to S from W at T and S does not do A at T' in W" and p is true in W"; and
- (3) for all worlds W" such that W" is accessible to S from W at T and S does not do A at T' in W" and p is true in W", there is a world W' such that:
  - (a) W' is accessible to S from W at T,
  - (b) S does A at T' in W',
  - (c) p is true in W', and
  - (d) the value for S at T of W' is greater than the value for S at T of W".

Roughly: S ought to do A, on the condition that p is true, iff S can do A while p is true, S can refrain from doing A while p is true, and any way of refraining from doing A-with-p is worse than some way of doing A-with-p. It is important to note that a conditional obligation, so construed, does not warrant the unrestricted detachment of an unconditional obligation when the condition in question is satisfied. Suppose that Matt can at T1 either attend a meeting in Los Angeles at T4, or attend a meeting in Chicago at T4, but not both, and that, in order to do the first, he must catch a plane at T2, while, in order to do the second, he must catch a plane at T3. Suppose also that he ought (absolutely) at T1 to attend the meeting in Los Angeles at T4, but also that he ought at T1 to attend the meeting in Chicago at T4 if he does not attend the meeting in Los Angeles at T4. Suppose, finally, that Matt, out of sheer will not attend the meeting in Los Angeles at T4.[7] It hardly follows that what he ought at T1 to do is attend the meeting in Chicago at T4, and (II), appropriately, does not imply that this does follow.

On the other hand, (II), again appropriately, does imply that certain restricted versions of detachment are valid. I have in mind two in particular, which may be called necessity detachment and deontic detachment. The former may be understood as follows: if S ought, at T in W, to do A at T', on the condition that p is true, and S cannot, at T in W, so act that p is false, then S ought, at T in W, to do A at T'. This holds because, if S cannot so act that p is false, then all worlds accessible to him are p-worlds; and so, if the best p-worlds accessible to him are A-worlds, then the best worlds accessible to him are A-worlds. Deontic detachment may be understood as follows: if S ought, at T in W, to do A at T', on the condition that p is true, and S ought, at T in W, so to act that p is true, then S ought, at T in W, to do A at T'. This holds because, if S ought so to act that p is true, then the best worlds

accessible to him are p-worlds; and so, if the best p-worlds accessible to him are A-worlds, then the best worlds accessible to him are A-worlds.

A final point to note about conditional obligation is its defeasibility. For example, while Matt ought at T1 to attend the meeting in Chicago at T4 if he doesn't attend the meeting in Los Angeles at T4, it is not the case that he ought at T1 to attend the meeting in Chicago at T4 if both he doesn't attend the meeting in Los Angeles at T4 and he commits suicide at T3.

The question now before us is this: how can prima facie obligation be understood in a manner that exploits the analysis of absolute obligation that has just been given?

2. PR1MA FACIE OBL1GAT1ON

To many, prima facie obligation has seemed to be obligation that is essentially subject to some sort of condition. Ross himself suggests the term "conditional duty" as an alternative to the term "prima facie duty".[8] And this may seem quite natural. Consider, for instance, the claim that one ought to keep one's promises, where this is taken to express a prima facie obligation. We may express this as: if one makes a promise to do A, then one ought to do A. How is this is to be analyzed?

2.1. Defective Analyses

It is clear that the following won't do as an analysis: that one makes a promise to do A implies (either materially, or strictly, or subjunctively) that one ought absolutely to do A. This is because such implication warrants the unrestricted detachment of obligation, whereas it is clear that the fact that one has made or will make a promise cannot by itself transform a prima facie obligation to keep it into an absolute obligation to do so.[9] An additional problem with the suggestion is that neither material nor strict conditionals are defeasible, and this would appear difficult to reconcile with the overridability of prima facie obligation. While one's prima facie obligation to keep a promise may on occasion also be an absolute obligation, it needn't be; and it won't be, even though one has made a promise, if something more pressing requires one's attention.

We might try analyzing the prima facie obligation to keep a promise to do A as follows: one ought absolutely either not to make a promise to do A or to do A. (The idea here is to have the absolute obligation apply to the entire conditional, itself reformulated as a disjunctive statement, and not just to its consequent).[10] But again this falls afoul of defeasibility. Suppose that one makes a promise at T1 to do A at T3 and that one therefore has a prima facie obligation at T2 to do A at T3. According to the proposal, then, one ought absolutely at T2 either not to make the promise at T1 or to do A at T3. But one cannot at T2 not make the promise at T1. Given the principle that, if one ought absolutely either to do X or to do Y and one cannot do X, then one ought absolutely to do Y, it follows that one ought absolutely at T2 to do A at T3. But this may not be the case; for, even though one has a prima facie obligation at T2 to keep one's promise and do A at T3, it may be that one has no absolute obligation at T2 to do so, because of some more pressing, contrary prima facie obligation.[11]

The fact that conditional absolute obligation, as analyzed in (II), does not warrant unrestricted detachment and is defeasible may suggest that it is well-suited to an analysis of the concept of prima facie obligation. But as Fred Feldman, one of the leading proponents of maximization, has noted, since necessity detachment is valid for statements of conditional absolute obligation, we cannot analyze the prima facie obligation to keep a promise to do A as a conditional absolute obligation to do A if one promises to do A. For, where it is also true that one cannot so act that one does not promise to do A (the most obvious type of case being one where the promise has already been made), we may detach from such a conditional absolute obligation the unconditional absolute obligation to do A. But, as already noted, we don't wish to say that one has an absolute obligation to do A just because one has a prima facie obligation to keep one's promises and one has already promised to do A.[12]

In light of the fact that the concept of prima facie obligation resists any of the foregoing analyses, Feldman concludes that it is merely a "mirage".[13] But this appears premature, as I shall now seek to demonstrate.

2.2. "All Else Being Equal"

The basic idea underlying the analysis that I shall propose is this. When one has a prima facie obligation to do A, doing A is of greater value, taken in an of itself, than not doing A (where "value" is used neutrally, as before in the account of absolute obligation, so as to be compatible with a wide variety of substantive theories of obligation). Another, and common, way of putting this is to say that doing A is better, all else being equal, than not doing A.[14] (Strictly, this characterization applies only to direct, and not to indirect, prima facie obligations, as will be explained shortly.)

This way of putting matters is quite revealing, for it indicates that prima facie obligation is indeed essentially absolute obligation conditioned in some way (and thus that Ross's alternative term "conditional duty" is an apt one). Thus it is sometimes said: one ought (absolutely), all else being equal, to act beneficently; one ought (absolutely), all else being equal, to keep one's promises. Nonetheless, the fact that prima facie obligation is absolute obligation conditioned in some way should not blind us to another fact that is frequently overlooked, and that is that prima facie obligation may itself be either unconditional or conditional. The prima facie obligation of fidelity is to be construed as a conditional prima facie obligation: if one makes a promise (or some other sort of commitment) to do A, then one has a prima facie obligation to do A. There is, it would seem, no call for fidelity tout court; it is fidelity to a commitment that is required, and, if there is no unconditional requirement to make a commitment (as in general there would seem not to be), so there is no unconditional requirement of fidelity. That is, while commitment-plus-fidelity is better, all else being equal, than commitment-plus-nonfidelity, we should not say that fidelity is better, all else being equal, than nonfidelity. For there are two ways to display nonfidelity — either by way of commitment-plus-nonfidelity, or by way of noncommitment-plus-nonfidelity — and only the former is tantamount to infidelity; there seems to be no reason to think that commitment-plus-fidelity is better, all else being equal, than noncommitment-plus-nonfidelity.

Having taken note of the possibility that prima facie obligation be either conditional or unconditional, we should note also that the detachment of an unconditional prima facie obligation from a conditional one is analogous to the detachment of an unconditional absolute obligation from a conditional one. That is, while unrestricted detachment is not valid, both necessity detachment and deontic detachment are valid. For example, even if one is going to make a promise to do A, one does not yet have for this reason alone an unconditional prima facie obligation to do A; rather, one simply has a conditional prima facie obligation to do A if one promises to do so. But if the promise is inevitable (for example, one has already made it), then one has an unconditional prima facie obligation to keep it; or if one has an unconditional prima facie obligation to make the promise, then one has an unconditional prima facie obligation to keep it.

Again, however, whether a prima facie obligation is itself unconditional or conditional, if it is properly characterized as an absolute obligation "all else being equal", then it is essentially an absolute obligation that is conditioned in some way; for, whether or not it is itself conditional, it carries with it this extra condition of "all else being equal". How are we to understand this extra condition?

It would make matters very simple if this condition were some sort of constant (call it "E" for "equal"). For then the obligation at issue would be straightforward conditional absolute obligation of the sort analyzed in (II) and both necessity detachment and deontic detachment would be clearly valid. Consider the claim that Faith has a fidelity-based prima facie obligation to pay Pam $10 ("P"), given her commitment ("C") to do so. This would, on the present suggestion, be understood to mean that she has an absolute obligation to do P on the condition that both C and E hold. Now consider necessity detachment, where this involves C's being inevitable. If Faith cannot so act that C does not hold, then all worlds accessible to her are C-worlds; and so, if the best (C&E)-worlds accessible to her are P-worlds, then the best E-worlds accessible to her are P-worlds. As for deontic detachment (where Faith has an unconditional prima facie obligation to do C): if the best (C&E)-worlds accessible to her are P-worlds and the best E-worlds accessible to her are C-worlds, then, again, the best E-worlds accessible to her are P-worlds.

What makes matters awkward is that it is clear that "all else being equal" is not to be interpreted as a constant. On the contrary, what it means varies from context to context; for it is a sort of indexical, in that "else" means "other than this". My suggestion is this: to say that S ought absolutely to do A, all else being equal, is to say, roughly, that any world, W, in which S does A is better than any other world, W', which is "just like" W except that in W' S does not do A. (As before, this in fact applies only to direct, and not to indirect, prima facie obligation. Again, more on this in a moment). Of course, this cannot be precisely correct, for no worlds can differ from one another with respect to just one such fact. If S does A in W but not in W', then any proposition which is distinct from but implies the proposition that S does A will also fail to be true in W'. What we need to say, then, is that W is as "close" to W as possible, given that in it S does not do A; that is, W' is "minimally different" from W with respect to S doing A.

This is not an unfamiliar notion. It, or something close to it, has been put to good use by several philosophers in the analysis of counterfactuals. To adapt what John Pollock has said:[15] the basic idea is that W' is a world that is obtained from W by making minimal changes which suffice to make it the case that S does not do A; truths in W that are in the appropriate sense "irrelevant" to or "independent" of S's doing A must also be true in W'; gratuitous changes are disallowed. I would like to be able to provide a precise analysis of this crucial concept, but I do not know how to do so. It might seem that all that needs to be said is this: W' is minimally different from W with respect to S doing A iff S does A in W but not in W' and every proposition that is true in W but which does not imply that S does A is also true in W'. But this won't do. Consider any two propositions p and q, neither of which alone implies that S does A but which jointly imply this. Given the current proposal, each of p and q will be true in W'; but then it will be true after all that S does A in W'. Pollock has himself proposed an analysis of minimal difference that does not succumb to this problem.[16] It may be that this analysis is acceptable, but I am reluctant to commit myself to it; for it is very complex and, besides, rests on the controversial notion of a "simple" proposition. Moreover, the analysis (as is typical in the treatment of counterfactuals) involves a past-future asymmetry, so that similarity of the past is weighted more heavily than similarity of the future; and I do not wish to presuppose any such asymmetry here. Formally, then, I shall invoke the concept of the closeness of possible worlds unanalyzed. I hope that my use of this concept will be sufficiently intuitive to render what I have to say about prima facie obligation both plausible and helpful, just as its use (or the use of some closely related concept) has proven helpful in the analysis of counterfactuals.

For reasons (having to do with detachment) that will become apparent shortly, I shall need to talk of those worlds that are minimally different for S at T in W with respect to S doing A; that is, the focus of attention will be on those worlds that are accessible to S. The canonical phrase- form in this context is therefore this: W" is minimally different from W', for S at T in W, with respect to S doing A at T'. This may be understood in terms of the closeness of worlds as follows:

(III) W" is minimally different from W', for S at T in W, with respect to S doing A at T' iff
- (1) W' is accessible to S from W at T and S does A at T' in W';
- (2) W" is accessible to S from W at T and S does not do A at T' in W"; and
- (3) there is no world W"' such that
  - (a) W"' is accessible to S from W at T,
  - (b) S does not do A at T' in W"', and
  - (c) W"' is closer to W' than W" is.

I shall also make use of the related concept of conditional minimal difference, which may be understood as follow:

(IV) W" is minimally different from W', for S at T in W, with respect to S doing A at T', on the condition that p is true, iff
- (1) W' is accessible to S from W at T and S does A at T' in W' and p is true in W';
- (3) W" is accessible to S from W at T and S does not do A at T' in W" and p is true in W"; and
- (3) there is no world W"' such that

(a) W''' is accessible to S from W at T,
(b) S does not do A at T' in W''',
(c) p is true in W1'', and
(d) W''' is closer to W' than W'' is.

Given this understanding of minimal difference, both unconditional and conditional, we may proceed to an analysis of prima facie obligation. I propose initially that we say this (revision will be called for later):

(V) S has, at Tin W, a direct unconditional prima facie obligation to do A at T' iff
(1) there are a world W' and a world W'' such that W'' is minimally different from W', for S at Tin W, with respect to S doing A at T'; and
(2) for all worlds W' and W'', if
(a) W'' is minimally different from W', for S at T in W, with respect to S doing A at T', then
(b) the value for S at T of W' is greater than the value for S at T of W''.

Similarly, we may say:

(VI) S has, at T in W, a direct conditional prima facie obligation to do A at T', on the condition that p is true, iff
(1) there are a world W' and a world W'' such that W'' is minimally different from W', for S at Tin W, with respect to S doing A at T', on the condition that p is true; and
(2) for all worlds W' and W'', if
(a) W'' is minimally different from W', for S at T in W, with respect to S doing A at T', on the condition that p is true, then
(b) the value for S at T of W' is greater than the value for S at T of W''.

Four points should be noted here.

First, instead of talking of prima facie obligations, we could talk of reasons; for there is a sense of "reason" according to which one has a reason to do something just in case one has a prima facie obligation to do it. Note that reasons may be actual or merely potential; I take this to mean that they may be unconditional or merely conditional.

Second, (V) and (VI) imply that, if S has a direct prima facie obligation to do A, then S can do A and S can refrain from doing A. A is, in other words, personally optional for S. This is in keeping with the analysis of absolute obligation provided in (I) and (II) above.

Third, (V) and (VI) are restricted to what I have called direct prima facie obligations. (Ross uses the term "direct" similarly).[17] If Ben has a direct prima facie obligation to act beneficently, this consists in every accessible beneficence-world being better than any accessible world that is minimally different with respect to beneficence. Similarly, if Faith has a direct prima facie obligation to pay Pam $10, if she so promises, this consists in every accessible promise-plus-payment world being better than any accessible promise-world that is minimally different with respect to payment. But not all prima facie obligations are direct. Some are indirect, and there are two ways in which they may be so. First, some may be called (following Ross again)[18] incidental. For example, if Ben cannot act beneficently without driving to Boston, then he has an incidental unconditional prima facie obligation to drive to Boston. Note that, since it would be highly surprising if all accessible worlds in which Ben drives to Boston were better than all accessible worlds that are minimally different in this respect, we have no reason to think that Ben has a direct unconditional obligation to drive to Boston. But, as long as he must drive to Boston to display beneficence, Ben does have an incidental, beneficence-based obligation to drive to Boston. Similarly, if Faith cannot keep her promise without going to the bank and withdrawing some cash, then she has an incidental conditional prima facie obligation to go to the bank and withdraw some cash. The principle here is straightforward. It may be roughly stated as follows: if S has a direct prima facie obligation to

do A and cannot do A without doing B, then S has an incidental prima facie obligation to do B.[19] More precisely, let us say, first of all:

> (VII) p requires q, for S at T in W, iff for all worlds W' accessible to S from W at T, if p is true in W', then q is true in W'.

We may then account for incidental prima facie obligations as follows:

(VIII) S has, at T in W, an incidental unconditional prima facie obligation to do B at T" iff for some act A and time T',
(1) S has, at T in W, a direct unconditional prima facie obligation to do A at T';
(2) S doing A at T' requires S doing B at T", for S at T in W;
(3) S doing A at T' is distinct from S doing B at T"; and
(4) there is a world W' such that W' is accessible to S from W at T and S does not do B at T" in W'.

(IX) S has, at T in W, an incidental conditional prima facie obligation to do B at T", on the condition that p is true, iff for some act A and T',
(1) S has, at T in W, a direct conditional prima facie obligation to do A at T', on the condition that p is true;
(2) S doing A at T' requires S doing B at T", for S at T in W;
(3) S doing A at T' is distinct from S doing B at T"; and
(4) there is a world W' such that W' is accessible to S from W at T and S does not do B at T" in W'.

The third clause in each is required to ensure that not all direct prima facie obligations are also said to be incidental. The fourth clause is required to preserve the thesis that whatever one has a prima facie obligation to do is personally optional.

The second way in which prima facie obligations can be indirect may be illustrated as follows. Suppose that Ben has an incidental prima facie obligation (grounded in beneficence) to drive to Boston and also has an incidental prima facie obligation (grounded in nonmaleficence) not to go to Boston without paying a visit to his mother, who lives there (in that, while it would not be mean of him not to visit his mother, it would be mean of him to go to Boston and not visit her). Then Ben has what I shall call a compound prima facie obligation (grounded in beneficence and nonmaleficence) to visit his mother. The principle at issue here may be roughly stated as follows: if S has a prima facie obligation (whether direct or incidental or, indeed, compound) to do A and also has a prima facie obligation not to do A without doing B, then S has a compound prima facie obligation to do B. More precisely:

(X) S has, at T in W, a compound unconditional prima facie obligation to do B at T" iff for some act A and T',
(1) S has, at T in W, a prima facie obligation to do A at T';
(2) S has, at T in W, a prima facie obligation not to both do A at T' and not do B at T"; and
(3) there is a world W' such that W' is accessible to S from W at T and S does not do B at T" in W'.

(XI) S has, at T in W, a compound conditional prima facie obligation to do B at T", on the condition that p is true, iff for some act A and time T',
(1) S has, at T in W, a prima facie obligation to do A at T', on the condition that p is true;
(2) S has, at T in W, a prima facie obligation not to both do A at T' and not do B at T"; and
(3) there is a world W' such that W' is accessible to S from W at T and S does not do B at T" in w .[20]

2.3. Three Objections

The foregoing account of prima facie obligation may appear objectionable for at least three reasons.

First, it might be thought that the account implies that whatever is directly prima facie obligatory is also absolutely obligatory (or at least absolutely permissible), for the following reason. According to the account, S

has a direct prima facie obligation to do A just in case, for every two accessible worlds that are just alike except in one S does A and in the other S doesn't do A, the former is better than the latter. But this appears to imply that the best accessible world (or worlds), W, is one in which S does A; for otherwise that accessible world that is just like W except in it S does A would be better than W.

Clearly such an implication would be unacceptable, but fortunately the present account doesn't have it. The term "just alike" is here not just rough but inaccurate. (Note that, when giving my rough, intuitive account of "all else being equal" above, I used the term "just like" and not the term "just alike".) In particular, one cannot assume that, if W" is minimally different from W' (for S at T in W) with respect to S doing A at T', then W' is minimally different from W" with respect to S not doing A at T'. Suppose, for example, that Ralph is faced with conflicting direct prima facie obligations of reparation and beneficence. Let us say that the best accessible world is one in which he satisfies the former but not the latter (he cannot satisfy them both). Next best would be his satisfying the latter but not the former. Worse still would be his satisfying neither. This may be depicted as follows (where "R" stands for "reparation" and "B" for beneficence"):

W_1: R & ~ B & …
W_2: ~ R & B & …
W_3: ~R & ~B & …

Given a proper filling-in of "... ", we can say that W2 is minimally different from W1 with respect to Ralph's not acting beneficently; and yet W3 (and not W1) is minimally different from W2 with respect to Ralph's acting beneficently. Thus what is prima facie obligatory need not also be absolutely obligatory or permissible.

A second objection stems from consideration of cases such as the following. Suppose that Gertrude owes Lucy a debt of gratitude. The most appropriate way for her to display gratitude would be to send Lucy lilies; daisies would be a poor second. But even though sending daisies would be inferior, it might seem that, insofar as it would constitute some display of gratitude, Gertrude has some reason, and hence a prima facie obligation, to send daisies; it's just that this obligation is overridden by the obligation to send lilies. But my account does not warrant our saying this. Presumably Gertrude has no direct prima facie obligation to send daisies, and doing so is not necessary for a display of gratitude and hence is not even incidentally prima facie obligatory (barring its being necessary for the satisfaction of some other direct prima facie obligation). In fact, it might appear that I cannot even say that Gertrude has an incidental prima facie obligation to send lilies, since even this appears not to be necessary for a display of gratitude (in that sending daisies would be an alternative, albeit inferior).

What I think we should say here is this. Granting that gratitude comes in degrees, let us say that "full" gratitude is gratitude to degree 10. Gertrude's sending lilies would thus constitute a display of gratitude to degree 10; her sending daisies would, let us say, only constitute a display of gratitude to degree 5. Now, while it is plausible to say that a display of gratitude is better, all else being equal, than a lack of a display of gratitude, it doesn't seem plausible to say that a display of gratitude to degree x is better, all else being equal, than a lack of a display of gratitude to degree x; for what if the lack of a display of gratitude to degree x is due to a display of gratitude to a greater degree y? Still, it is plausible to say that a display of gratitude to at least degree x is better, all else being equal, than a lack of a display of gratitude to at least degree x. This observation allows us to say that Gertrude does indeed have an incidental obligation to send Lucy lilies (given that there is no degree of gratitude greater than 10 and that sending lilies is the only way to display gratitude to that degree); and, while it won't allow us to say that she also has an incidental obligation to send Lucy daisies, it does allow us to say that she has an incidental obligation either to send lilies or to send daisies (given that nothing else would constitute a display of gratitude to at least degree 5). And this indeed seems to be what should be said.

In fact, the matter is even more complex than this. Suppose that Gertrude ought, from gratitude, to send Lucy lilies, and also ought, from gratitude, to send Chrissie chrysanthemums, but cannot do both. Here, it seems, gratitude requires the impossible, and yet my account implies that whatever is prima facie obligatory is possible. The solution, of course, is that gratitude, more fully characterized, is gratitude-to-a-benefactor (or, perhaps,

would-be benefactor, or presumed benefactor), just as fidelity is fidelity-to-a-promisee (or someone to whom another sort of commitment has been made). (If we talk of "what gratitude requires" without qualification, then, we must mean what gratitude to every benefactor requires; so too for fidelity, reparation, and any other obligation which is owed to a particular person or persons.)[21] My account implies, properly, that Gertrude can satisfy each of her prima facie obligations of gratitude-to-Lucy and gratitude-to-Chrissie; it does not imply, nor should it, that she can satisfy them both.

A third objection is more troublesome. It is that the account implies that certain direct prima facie obligations that are intuitively thought of as being merely conditional turn out also to be unconditional. Consider fidelity. This has been assumed (for purposes of illustration) to constitute a direct conditional prima facie obligation, in that all commitment-plus-fidelity-worlds have been assumed to be better than any worlds that are minimally different with respect to fidelity. But notice that one cannot (it seems) display fidelity in the absence of a commitment; that is, all accessible fidelity-worlds are also commitment-worlds. But then, if all accessible commitment-plus-fidelity-worlds are better than their closest commitment-plus-nonfidelity-worlds, it must also be true that all accessible fidelity-worlds are better than their closest nonfidelity worlds. In other words, the direct prima facie obligation of fidelity will be not just conditional but unconditional. The same applies to other (alleged) conditional direct prima facie obligations such as those of gratitude and reparation. For one cannot (it seems) display gratitude in the absence of that condition (namely, being a beneficiary) on which it is conditionally prima facie obligatory; likewise, one cannot make reparation for a nonexistent harm.

It is tempting at first simply to bite the bullet here but then go on to point out that there are still certain intimately related direct prima facie obligations that are merely conditional. For example, even if Faith cannot display fidelity to her promise to pay Pam $10 without making that promise, still she can pay Pam $10 without promising to do so. Thus, while all accessible promise-plus-payment-worlds may be better than their closest accessible promise-plus-nonpayment-worlds, there is no need to think that all accessible payment-worlds are better than their closest accessible nonpayment-worlds. And isn't that what really matters here?

Well, no, it isn't. There are two problems with this. First, on some occasions, it may not be possible to identify a conditionally obligatory act in such a way that the condition is not included in the identification. While Faith can do what fidelity requires (namely, pay Pam $10) without promising to do it, it may be that Ralph cannot do what reparation requires (mending a broken vase, say) without committing the relevant harm. Secondly, if it is admitted that Faith has a direct unconditional prima facie obligation to display fidelity (even if it is also pointed out that her direct prima facie obligation to pay Pam $10 is merely conditional), it must then also be admitted, given the foregoing account, that she has an incidental prima facie obligation to make the relevant promise that is unconditional. Similarly, if it is admitted that Ralph has a direct unconditional prima facie obligation to make some reparation, then it must also be admitted that he has an incidental prima facie obligation to commit the relevant harm that is unconditional. This is surely to admit too much. The original account must therefore be revised.

We must deny not only that Faith has an unconditional prima facie obligation to pay Pam $10 but also that she has an unconditional prima facie obligation to display fidelity to the commitment to pay Pam $10 (given that she has not already made this commitment and that there is no independent unconditional reason to make it). (Similar remarks pertain to the prima facie obligations of gratitude and reparation.) What I think warrants such denial is an observation already made above: while commitment-plus-fidelity is better, all else being equal, than commitment-plus-nonfidelity, it is not better, all else being equal, than noncommitment-plus-nonfidelity; for nonfidelity constitutes infidelity only when there is an actual commitment. It is in virtue of this fact that the prima facie obligation of fidelity is merely conditional. But how can this be? After all, consider a situation where it is possible to display a great deal of maleficence to one person while also displaying minor beneficence to another. We may suppose in this case that maleficence-plus-beneficence is better, all else being equal, than maleficence-plus-nonbeneficence, and also that maleficence-plus-beneficence is not better, all else being equal, than nonmaleficence-plus-nonbeneficence. Yet this doesn't make the prima facie obligation of beneficence a

merely conditional one. So why should the corresponding fact about commitment-plus-fidelity make the prima facie obligation of fidelity a merely conditional one?

The answer is to be found in the fact that commitment is a prerequisite of fidelity, while maleficence is not a prerequisite of beneficence. This is the very feature about fidelity that has given rise to the present problem, and it points to the proper solution. Commitment is a prerequisite of fidelity, not just in the sense that there cannot be fidelity without a commitment, but in the sense that fidelity requires that a commitment precede it. Commitment and fidelity thus come in a well-ordered "package", as it were, and the package must commend itself ab initio if fidelity is to be unconditionally obligatory. (Similar remarks pertain to gratitude and reparation and their respective initial conditions.) To put this point more precisely, let us first say this:

(XII) W" is minimally different from W', for S at T in W, with respect to S doing A at T', on the condition that p varies, iff
- (1) W' is accessible to S from W at T and S does A at T' in W' and p is true in W';
- (2) W" is accessible to S from W at T and S does not do A at T' in W" and p is not true in W"; and
- (3) there is no world W"' such that
  - (a) W"' is accessible to S from W at T,
  - (b) S does not do A at T' in W"',
  - (c) p is not true in W"', and
  - (d) W"' is closer to W' than W" is.

With this in hand, we may now revise (V) by appending an extra clause as follows:

(V')S has, at T in W, a direct unconditional prima facie obligation to do A at T' iff
- (1) there are a world W' and a world W" such that W" is minimally different from W', for S at T in W, with respect to S doing A at T';
- (2) for all worlds W' and W", if
  - (a) W" is minimally different from W', for S at T in W, with respect to S doing A at T', then
  - (b) the value for S at T of W' is greater than the value for S at T of W11; and
- (3) for all propositions p, if
  - (a) S has, at T in W, a direct conditional prima facie obligation to do A at T', on the condition that p is true,
  - (b) for some event E and time T" earlier than p is the proposition that E occurs at T",
  - (c) S doing A at T' requires p, for S at T in W, and
  - (d) there is a world accessible to S from W at T in which p is not true, then
  - (e) for all worlds W' and W", if
    - (i) W" is minimally different from W', for S at T in W, with respect to S doing A at T', on the condition that p varies, then
    - (ii) the value for S at T of W' is greater than the value for S at T of W".

The "official" account of prima facie obligation is thus as follows: replace (V) by (V') and leave (VI) (which, indeed, is presupposed by (V')) and (VII) — (IX) as they are.

Finally, it should be noted that this account accommodates the possibility that the value at issue be significantly agent- or time-relative, inasmuch as the relevant value of the worlds is the value they have for S at T. Let me now elaborate on this with respect to Ross's theory. I said earlier that at least some of the values that ground prima facie obligations are thought by Ross to be significantly agent-relative. Fidelity, for example, is clearly a value of this sort. Ross says that each agent always has a prima facie obligation to see to it that he keeps his promises.[22] I take this to imply that each of us always has a (conditional) prima facie obligation to see to it that we ourselves display fidelity. Now it might be that Ross also believes that each of us always has a prima facie obligation to see to it that others display fidelity; for he might believe that such displays are intrinsically valuable, and he does say that we do always have a prima facie obligation to do that which will promote such

value.[23] But even if this is so, it does not detract from the fact that it seems pretty clear that, according to Ross, for each of us a special value attaches to our displaying fidelity. So too for gratitude and reparation, it seems.[24] Whether or not this is so for other prima facie obligations that Ross lists is unclear to me. Does he believe, for instance, that there is a special value for each of us to be attached to our not harming innocent people, even if there is also a value to be attached to preventing others from harming innocent people? I am not sure, although certain of his remarks suggest that he does believe this.[25]

2.4. Detachment

The move from (V) to (V') blocks unrestricted detachment. Even if Faith will promise to pay Pam $10, the prima facie obligation to display fidelity to this commitment is merely conditional; for clause (3) of (V') is not satisfied. To be acceptable, though, the account must not only block unrestricted detachment, it must also validate necessity detachment and deontic detachment. And so it does. Consider necessity detachment first. If Faith cannot avoid the commitment to pay Pam $10, then clause (3) of (V') is satisfied after all, and trivially so; for there will then be no accessible worlds that are minimally different with respect to Faith's displaying fidelity to this commitment, on the condition that her making this commitment varies. Thus her conditional prima facie obligation to display fidelity is transformed into an unconditional prima facie obligation to do so, and her obligation to pay Pam $10 is likewise transformed.

Deontic detachment also holds, although the story here is more complicated. Let us look first at deontic detachment having to do with direct prima facie obligations. Suppose that Faith has already made a promise (promise 1) to make another promise (promise 2) to pay Pam $10. Then, by necessity detachment, she has an unconditional prima facie obligation to make promise 2. Given the conditional prima facie obligation to pay Pam $10 if she makes promise 2, it should turn out that Faith has an unconditional prima facie obligation to pay Pam $10. And it does. Consider these world-schemata (where "P2" stands for "promise 2" and "P" for "payment"):

W_1: P2 & P & …
W_2: P2 & ~P & …
W_3: ~ P2 & P & …

We already know (given a proper filling-in of "... ") that, because of the direct unconditional prima facie obligation to make promise 2, W1 is better than W3. We also know that W1 is better than W2, for this is what the direct conditional prima facie obligation to pay Pam $10 if promise 2 is made amounts to. Thus (P2 & P) is better, all else being equal, than ,,,(P2 & P), and so Faith has a direct unconditional prima facie obligation to do (P2 & P).[26] Thus she also has an incidental prima facie obligation to do P.

Consider, next, incidental prima facie obligations. Suppose that Faith's prima facie obligation to make promise 2 is not grounded in any already-made promise but rather in beneficence. Beneficence requires that she make promise 2; her obligation to do so is thus incidental and not direct. Consider these world-schemata (where "B" stands for "beneficence" and "P" for "payment"):

W_1: B & P & ...
W_2: B & ~ P & ...
W_3: ~B & P & ...

Since beneficence is directly unconditionally obligatory, W1 is better than W3. In addition, since B requires P2 (where "P2" stands for "promise 2"), W1 and W2 must have the following forms, respectively:

W_1: B & P_2 & P & …
W_2: B & P_2 & ~ P & …

Since P is directly prima facie obligatory on the condition that P2 occurs, W1 is better than W2. Hence (B & P) is unconditionally obligatory, and so P is as well.

Consider, finally, compound prima facie obligations. Suppose that beneficence requires that Faith visit her mother and that nonmaleficence requires that she not visit her mother without making promise 2. Then she has an unconditional compound obligation to make promise 2. Consider these world-schemata (where "B" stands for "beneficence", "N" for "nonmaleficence", and "P" for "payment"):

W_1: B & N & P & …
W_2: B & N & ~ P & …
W_3: B & ~ N & P & …
W_4: ~ B & N & P & …

Given the direct unconditional prima facie obligation of beneficence, we may say that W1 is better than W4. Given the direct unconditional prima facie obligation of nonmaleficence, we may say W1 is better than W3. In addition, since (B & N) requires P2 (where "P2" stands for "promise 2"), W1 and W2 must have the following forms, respectively:

W_1: B & N & P_2 & P & …
W_2: B & N & P_2 & ~ P & …

Since P is prima facie obligatory on the condition that P2 occurs, W1 is better than W2. Hence (B & N & P) is unconditionally obligatory, and so P is as well.

2.5. Overriding

What it is for one prima facie obligation to override another may now be accounted for in a straightforward manner. Roughly, S's prima facie obligation to do A overrides his prima facie obligation to do B if every accessible world in which he does B is worse than some accessible world in which he does A. More precisely:

(XIII) S's prima facie obligation, at T in W, to do A at T' overrides S's prima facie obligation, at T in W, to do B at T' iff
- (1) S has a prima facie obligation, at T in W, to do A at T';
- (2) S has a prima facie obligation, at T in W, to do B at T'; and
- (3) for all worlds W" such that W" is accessible to S from W at T and S does B at T' in W", there is a world W' such that:
  - (a) W' is accessible to S from W at T,
  - (b) S does A at T' in W', and
  - (c) the value for S at T of W' is greater than the value for S at T of W".

This applies to all prima facie obligations, whether unconditional or conditional, direct or indirect.

Notice that it is possible that direct obligations override in a manner that fails to match the manner in which their corresponding incidental obligations override. Suppose that Ben has a direct prima facie obligation of beneficence and that on a certain occasion this overrides a direct prima facie obligation of fidelity. Suppose also that he has the corresponding incidental obligations to drive to Boston and to drive to Baltimore, respectively. It may be that his obligation to drive to Boston does not override his obligation to drive to Baltimore; indeed, it may even be that the latter overrides the former. This would happen if there is yet another direct prima facie obligation, say an obligation of reparation, that overrides the obligation of beneficence and the fulfilment of which itself requires that Ben drive to Baltimore.

3. ROSS ON PRIMA FAC1E OBL1GATION

Since it was Ross who introduced the concept of prima facie obligation to us, I have relied fairly heavily on what he says in the presentation of my own analysis of the concept. I have invoked his distinction concerning direct and incidental obligations, and I have made free use of some of his examples of prima facie obligation (fidelity, beneficence, gratitude, reparation, and so on). It is appropriate at this point to try to determine to what extent my account squares with Ross's, at least with regard to certain key issues.

3.1. Conditional and Unconditional Prima Facie Obligation

First, it might be thought that Ross does not distinguish between conditional and unconditional prima facie obligation, inasmuch as he sometimes refers to all prima facie obligation by the term "conditional duty". But in fact I think that Ross does make this distinction but simply uses different terminology to do so. I have in mind what he says about what he calls "special" and "general" prima facie obligations.[27] Special obligations, he says, rest on the performance of some prior act (such as the making of a promise, the infliction of an injury, or the acceptance of a benefit); the obligations (of fidelity, reparation, and gratitude, respectively) are therefore conditional on the performance of these acts. General obligations (among which Ross counts the obligations of justice, beneficence, and nonmaleficence) do not rest on the performance of such prior acts and may therefore be said to be unconditional.

At least, this is how I understand Ross's distinction between special and general obligations and what he means by "rest on". Note that, in some other, more liberal sense of "rest on", all obligations, both conditional and unconditional, will rest on the satisfaction of some condition; for the performance of any act requires that certain conditions be satisfied. For example, just as no one can keep a promise unless he has made one and no one can make reparation for an injury unless one has been inflicted, so too no one can administer justice unless there is occasion for doing so and no one can act beneficently unless someone is available as a potential beneficiary. So what Ross means by "rest on" cannot be just this. Rather it is, I suggest, what I have indicated, namely, that, with respect to fidelity, not only does fidelity require commitment, but noncommitment-plus-nonfidelity is no worse than commitment-plus-fidelity, and so fidelity isn't unconditionally obligatory; similarly, noninfliction-of-harm-plus-nonreparation is no worse (indeed, it is better, given the prima facie duty of nonmaleficence) than infliction-of-harm-plus-reparation, and so reparation isn't unconditionally obligatory; and so on. But with respect to beneficence, for example, there is no "prior act" like that of commitment (or the infliction of harm) which it is prima facie permissible to forgo, and so nothing that is morally on a par with noncommitment-plus-nonfidelity. All nonbeneficence is morally on a par with infidelity (that is, with commitment-plus-nonfidelity). Or so I understand Ross. Of course, someone might take the view that beneficence is in general only conditionally required, in that its obligatoriness rests on the performance of some prior act (such as that of becoming a neighbor of the potential beneficiary, in some sense of "neighbor" according to which not everyone else is one's neighbor); but I am sure that this is not Ross's view.

3.2. The Self-evidence of Prima Facie Obligations

Ross says that the direct prima facie obligations he mentions are self- evident. His remarks concerning this matter clearly indicate that he believes that it is a necessary truth that certain acts are prima facie obligatory (whether conditionally or unconditionally).[28] Of course, such a claim goes well beyond what a mere conceptual analysis could be thought to imply, but is it consistent with what I have said?

This is a complicated matter. The account so far provides an analysis of what it is to claim that a particular agent, S, has a prima facie obligation at a particular time. Unless S exists necessarily, no such claim can be a necessary truth. Still, we can allow for necessary truths concerning prima facie obligation by expanding the account that has already been given. We can, in particular, say the following:[29]

(XIV) S has a permanent unconditional prima facie obligation in W to do A if for all times T and T', if
  (1) there is a world accessible to S from W at T in which S does A at T', and
  (2) there is a world accessible to S from W at T in which S does not do A at T', then
  (3) S has an unconditional prima facie obligation, at T in W, to do A at T'.

(XV) there is a universal unconditional prima facie obligation in W to do A if for all persons S, if

(1) S exists in W, then
(2) S has a permanent unconditional prima facie obligation in W to do A.

(XVI) there is a necessary unconditional prima facie obligation to do A if for all worlds W, there is a universal unconditional prima facie obligation in W to do A.

A similar account can be given for conditional prima facie obligation. One can then go on to claim, as Ross appears to do, that, if ever an agent has a prima facie obligation to do some act A, this is grounded in the fact that there is a necessary direct prima facie obligation to do either A itself or some other act for which A is necessary.[30] While the analysis that I've provided does not imply this claim, it is certainly consistent with it.

Nonetheless, there are complications, one minor, one major, in trying to reconcile what Ross has to say about the self-evidence of certain particular prima facie obligations with the account that I have provided. The minor complication concerns the (alleged) prima facie obligation of fidelity. While I have allowed myself to invoke this obligation for the sake of illustration above, it is in fact doubtful that there is a necessary direct prima facie obligation of fidelity of the sort that Ross alleges. For what he appears to say is that every commitment necessarily gives rise to a prima facie obligation to fulfil that commitment.[31] This cannot be right. For one thing, people sometimes promise to do something that they in fact cannot do. The account that I have given implies that there is no prima facie obligation to keep such promises.[32] Of course, even if this is what Ross says, there is an easy qualification available, and that is to say that every commitment necessarily gives rise to a prima facie obligation to fulfil it, if the agent can fulfil it. But there is reason to think that this, too, is false. For it often happens that people promise to do what is in fact wrong, absolutely. What reason is there then to think that all worlds accessible to the agent in which he displays fidelity are better than all accessible worlds that are minimally different in this respect? For any world that is minimally different in this respect will at least involve his not committing the absolute wrong in question.[33] Once again, though, a qualification is perhaps available, although it is not so easy to see just how it should go.

The major complication is this, and it applies to all the particular prima facie obligations that Ross alleges to be self-evident, except perhaps the prima facie obligation to promote intrinsic value. The account that I have provided relies on the notion of minimal difference, where this is explicitly tied to worlds that are accessible to the agent in question. But surely it is possible that some agents be such that they cannot avoid some evil's occurring if they satisfy some Ross-alleged direct prima facie obligation; in such a case, given that the evil is sufficiently great (so that it renders any world in which it occurs sufficiently bad, and any world in which it does not occur sufficiently good relative to any world in which it does occur), it will not be the case that all accessible worlds in which the alleged obligation is satisfied are better than all accessible worlds that are minimally different in this respect. Hence this alleged obligation does not hold for such people; hence it does not hold of necessity; hence it is not self-evident that it holds. As an illustration of this, consider the possibility that Martha will commit mass murder if and only if Norman acts nonmaleficently, and that Norman can do nothing about this. We may therefore suppose (barring certain Kantian considerations) that not all worlds accessible to Norman in which he acts nonmaleficently are better than all worlds accessible to him that are minimally different in this respect. In fact, it may well be that the very reverse is true, so that Norman has (on my account) a prima facie obligation to act maleficently.

Thus it cannot be, on the account that I have provided, that there is a necessary unconditional prima facie obligation to act nonmaleficently. So, too, for beneficence and other obligations that Ross calls "general". Similarly, the conditional obligations of fidelity, gratitude, and so on that Ross calls "special" cannot hold of necessity, on the account that I have provided.

If we are to accommodate what Ross says about the self-evidence of such prima facie obligations, therefore, a different account must be given. Let me just hint here at how I think this account should go. The account would not restrict our focus to worlds that are accessible to the agent. We would begin with this revision of (VI):

(VI') S has at T a direct conditional prima facie obligation to do A at T', on the condition that p is true, iff

(1) there are a world W and a world W' such that W' is minimally different from W with respect to S doing A at T', on the condition that p is true;
(2) for all worlds W and W', if
(a) W' is minimally different from W with respect to S doing A at T', on the condition that p is true, then
(b) the value for S at T of W is greater than the value for S at T of W'.

Changes to (V') and to (VII) — (XVI) would then be made accordingly, and in this way the present problem concerning self-evidence would be obviated.

Why, then, did I not simply proceed in this fashion in the first place? The answer has to do with detachment. It is not that I think such a procedure mistaken; on the contrary, it would, if carried out in full detail, give us an account of what might be called prima facie obligations in principle. We could then say that everyone necessarily has a reason, in principle, to act nonmaleficently; this would apply even to Norman. But the problem is that, while versions of necessity detachment and deontic detachment would indeed hold for conditional prima facie obligations in principle, we would not be able to infer, as we surely wish to do, that, for example, someone for whom a promise is inevitable has an unconditional prima facie obligation to keep it. For more than mere inevitability would now be required for necessity detachment; what would be required is that the condition in question hold in all worlds, not just in all worlds accessible to the agent. Similarly, for deontic detachment, the unconditional obligation that triggers the detachment would itself have to be an obligation in principle, and this would block certain detachments that we intuitively wish to accept (such as that of detaching an unconditional obligation to do A by appealing to an unconditional obligation to promise to do A, the latter having itself been detached — by necessity detachment — due to considerations having to do merely with inevitability and not with metaphysical necessity).

The fact is that there are, I think, two types of prima facie obligations. There are prima facie obligations in principle, and there are prima facie obligations in practice. The detailed account that I have provided is of prima facie obligations in practice; the sketch of an account that I have intimated in the penultimate paragraph is of prima facie obligations in principle. We need both. We need the latter to accommodate what Ross says about self-evidence and necessity; we need the former, I believe, to accommodate common intuitions concerning detachment.

3.3. Prima Facie Obligation and Absolute Obligation

Let me turn, finally, to what Ross has to say about the relation between prima facie obligation and absolute obligation. Here he appears to make two important claims. They are these (put somewhat roughly):

(XVII) S has a prima facie obligation to do A if S would have an absolute obligation to do A if any prima facie obligation that S has not to do A were overridden;[34]
(XVIII) S has an absolute obligation to do A if S has a prima facie obligation to do A that overrides any prima facie obligation that S has not to do A.[35]

It might be thought that any account of obligation that purports to capture Ross's concept of prima facie obligation must imply that both of these theses are true. But the fact is that neither thesis is implied by my account (although each is consistent with it). The implication fails because my account of prima facie obligation incorporates a certain generality ("all worlds in which S does A are better than all worlds that are minimally different in this respect"), and such generality, which I take to be essential to prima facie obligations (indeed essential, in some form or other, to reasons of all sorts), is not incorporated in my account of absolute obligation. Does this show up a failing in either of these accounts?

I don't think so. Ross seems to think that there are several basic, independent moral values (such as those of fidelity, gratitude, and so on); that nothing else is morally valuable (at least with respect to the determination of obligation) unless it is a means to the promotion of these values; and that there is always, indeed necessarily, a

reason to promote these values. If we grant all this, and we also grant (as seems natural) that the moral value of a world is a direct function of those moral values which are promoted at that world, then, given the foregoing accounts of prima facie and absolute obligation, it seems that we could argue cogently for both (XVII) and (XVIII); for there will be nothing that is morally valuable that does not correspond to some prima facie obligation.[36] But whether or not we should grant all this seems to me to go well beyond questions of conceptual analysis. This being the case, it is not a failure in the foregoing accounts that neither (XVII) nor (XVIII) is implied by them.

I have said that my account of obligation is consistent with (XVIII). This might, however, be disputed on the grounds that Ross conceives of (XVIII) as an analysis .[37] If absolute obligation is to be understood in terms of prima facie obligation, as this interpretation of (XVIII) contends, then it is not to be analyzed independently of the concept of prima facie obligation, as (I) contends.

Now, I think it must be granted that, if (XVIII) is accepted as an analysis of absolute obligation, my proposed analysis of absolute obligation must be rejected.[38] I therefore reject (XVIII) as an analysis (though I grant that it may express a necessary truth). How does Ross himself view (XVIII)? The evidence seems to me inconclusive: while some passages[39] in The Right and the Good may appear to suggest that he does accept (XVIII) as an analysis, he never states this explicitly; moreover, other passages[40] strongly suggest that he eschews any such analysis. At any rate, if Ross does accept (XVIII) as an analysis, I oppose him on this point. I leave it to the reader to judge whether this indicates an inadequacy in my account.

One reason to suspect that it does indicate an inadequacy might be this. As has been noted, Ross claims that (certain of) one's prima facie obligations are self-evident, but he also claims that one's absolute obligations are not and that one may appeal to the former when trying to determine the latter. This may appear to indicate a priority of prima facie obligation over absolute obligation that is inconsistent with (I). Not so. As Frank Snare has pointed out, the logical or conceptual independence of absolute obligation from prima facie obligation is not inconsistent with the former being epistemologically dependent on the latter.[41] Similarly, one might claim that sometimes it is appropriate to appeal to one's prima facie obligations in order to explain what one's absolute obligations are.[42] But, once again, the account given in (I) is not inconsistent with such priority of prima facie obligation over absolute obligation. For it may well be that certain matters are, and that we find them to be, of fundamental moral significance — well-being, perhaps, or respect, or integrity, or fidelity, or justice, or freedom, or something else — such that our absolute obligations turn on, and we believe them to turn on, whether these values are promoted. Far from wishing to deny this, I believe it to be true. No doubt, then, we will agree that there are prima facie obligations, both in principle and in practice, having to do with the promotion of these values, so that sometimes, when we deliberate about what we ought absolutely to do in some situation, we may bear in mind that these prima facie obligations apply in this situation. But this is quite consistent with the analyses of absolute and prima facie obligation that have been proposed here.[43]

NOTES

1. "Iff" is short for "if and only if". I here intend for formulae of the form "p iff q" to be understood as at least partial analyses, in that what is on the right side of "iff" is intended to reveal, at least in part, the conceptual structure of what is on the left side. This holds until I present formulae (XVII) and (XVIII) in Section 3 below, where "iff" is intended merely to signify a strict equivalence.
2. The main inspiration of the accounts given in (I) and (II) (below) is Fred Feldman's account of obligation, presented in [3], chs. 2 and 4. Clause (l) of (I) is in fact redundant, being implied by the conjunction of clauses (2) and (3), but its explicit statement helps to emphasize that "ought" is understood to imply "can". Also, if is possible that there be an infinite number of worlds accessible to S and that these "alternate" with respect to whether or not he does A in them, then (I) needs to be complicated by adding a clause to the effect that W' in clause (3) is such that no accessible world in which S does not do A is of greater value than it.
3. This treatment of acts is not, I think, essential to the sort of maximizing theory that is at issue here, but it does simplify the presentation of it.

4. Compare Feldman's account of obligation in [3], ch. 2, which constitutes a modified version of traditional consequentialism.
5. Cf. Vallentyne, [12], p. 120, including n. 10 on p. 124.
6. [13], p.98.
7. This case is borrowed from [14], p. 67.
8. [9], p. 19.
9. Although I do not have the space to go into the matter here, I would contend that this problem of detachment also vitiates two much more plausible analyses, proposed in [1] and [11], of the concept of prima facie obligation in terms of a subjunctive conditional.
10. Compare the analysis proposed by Hintikka in [4], pp. 91ff.
11. Further problems with Hintikka's sort of analysis are discussed in [3], p. 182, and [8].
12. [3], pp. 130-1.
13. [3], p. 144. By this he means that there really is no such concept; terms that appear to express it are in fact meaningless or express some other concept.
14. Cf. [6] for an extended investigation into this characterization of prima facie obligation.
15. [7], pp. 70-1.
16. [7], pp. 70ff.
17. [9] p.46.
18. [9], p.46.
19. This principle may seem to be vulnerable to a paradox like that of the Good Samaritan. I don't think that it is, but I will not pursue this here.
20. A still more precise account of compound prima facie obligation, both unconditional and conditional, would (in order to avoid any appearance of circularity) distinguish between level-1 compound obligations, which would themselves be generated by obligations that are not themselves compound, and other level-n compound obligations, which would be generated by level-(n--1) compound obligations.
21. Perhaps, for example, nonmaleficence is such an obligation while beneficence is not. This has to do with whether persons have a right to the performance of an obligatory act. I shall not venture further into this territory here.
22. [9], pp. 17ff.
23. On the obligation to promote intrinsic value, see [9], pp. 24 and 39. In fact, however, Ross's discussion in [9], Ch. 5, especially p. 140, suggests that he does not believe that each of us always has, in virtue of the prima facie obligation to promote intrinsic value, a prima facie obligation to see to it that others display fidelity. For there he lists four intrinsic goods — virtue, pleasure, the allocation of pleasure to the virtuous, and knowledge — and he explicitly says that he doubts the existence of any others (save combinations of those mentioned). It might be thought that Ross takes fidelity to be a virtue, but this would be a mistake. For he takes virtue to concern motive, and he explicitly divorces the duty of what he calls fidelity form motive. (See [9], p. 22 and also p. 5.) Still, Ross does say ([9], p. 36) that fidelity can be bonific — that is, instrumentally good — and so he clearly thinks that, on those occasions on which it is bonific, there is a prima facie obligation to promote it.
24. See [9], pp. 27,30, and 41-2.
25. See [9], p. 22.
26. Clause (3) of (V') is satisfied since the only relevant prerequisite of (P2 & P) is promise 1. Note that there is no need in this context to compare W_1 with W_4: $\sim P_2$ & $\sim$ P & ... For, although this world is a $\sim(P_2$ & P)-world, it is not as close to W_1 as either W_2 or W_3 is. (Still, I think that W_1 is better than W_4, and that this could be satisfactorily proven.)
27. [9], p.27.
28. [9], pp. 29-30 and 32-3.
29. It is here that my account, as presented, clearly presupposes that acts can be done more than once and by more than one person. But see note 3 again.
30. Or, we might add, some other acts which imply that one has a compound prima facie obligation to do A. Ross does not appear explicitly to acknowledge the possibility of compound prima facie obligations.
31. This is most evident at [9], pp. 37 and 40, but see also pp. 21, 24, 27 and 28.

32. Some philosophers would of course respond by denying that "ought" implies "can". (See, e.g., [10], pp. 118-9.) (While Ross endorses the thesis that "ought" implies "can" on p. 5 of [9], he is there taking "ought" to express absolute obligation. He says nothing about whether the thesis holds when "ought" is taken to express prima facie obligation. Pietroski, on the other hand, is not silent on this. While he explicitly accepts that "ought" implies "can" where "ought" expresses absolute obligation, he rejects this claim where "ought" expresses prima facie obligation. See [6], pp. 494-5 and 504-5.) I shall not discuss this here. Others would respond by claiming that apparent promises to do what is impossible are in fact not genuine promises, and then claim that every genuine promise does in fact ground a prima facie obligation to fulfil it.(See, e.g., [5].) This seems mistaken, in that it seems possible that a person should sincerely, and thus genuinely, promise to do something, in the mistaken belief that he can do that thing.
33. Compare what Donagan says about promising in [2], pp. 92-3.
34. [9], p. 19.
35. [9], p. 41. Concerning the "iff" in both (XVII) and (XVIII), see note 1 above.
36. This is not to say that it is easy to see just how this argument would go. There are a number of complications, having to do with the distinctions between conditional and unconditional obligations, obligations in principle and obligations in practice, permanent versus universal versus necessary obligations, and with other matters.
37. Presumably he does not conceive of (XVII) as such, since it would be circular.
38. This claim presupposes that there cannot be two distinct acceptable analyses of a single concept, where one of these analyses itself employs concepts that the other does not.
39. [9], pp. 46-7.
40. [9], pp. I0-2 and 16. On p. 46 Ross says that the ground of the absolute rightness of an act is that its prima facie rightness outweighs its prima facie wrongness, and on p. 10 he says that "the reason for an action's being [absolutely] right is evidently not the same thing as its [absolute] rightness".
41. [11] p 244.
42. Cf. again [9], p. 46, Cf. also [6], p. 513.
43. Ted Sider provided me with a series of detailed comments which were extraordinarily penetrating and helpful and for which I am extremely grateful. I am grateful, too, for help from Fred Feldman, Ishtiyaque Haji, Terry McConnell, Gary Rosenkrantz, and Peter Vallentyne.

REFERENCES

[1] Atwell, J.: 1978, 'Ross and Prima Facie Duties'. *Ethics*, 88,240-9.
[2] Donagan, A.: 1977, *The Theory of Morality*. Chicago: University of Chicago Press.
[3] Feldman, F.: 1986, *Doing the Best We Can*. Dordrecht: D. Reidel.
[4] Hintikka, J.: 1970, 'Some Main Problems in Deontic Logic'. In Risto Hilpinen (ed.) *Deontic Logic: Introductory and Systematic Readings*. Dordrecht: D. Reidel, 59-104.
[5] Martinich, A. P.: 1987, 'Obligation, Ability and Prima Facie Promising'. *Philosoph* 17, 323-30.
[6] Pietroski, P. M.: 1993, 'Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory'. *Ethics*, 103, 489-515.
[7] Pollock, J. L.: 1976, *Subjunctive Reasoning*. Dordrecht: D. Reidel.
[8] Purtill, R. L.: 1973, `Deontically Perfect Worlds and Prima Facie Obligations'. *Philosophia*, 3, 429-38.
[9] Ross, W. D.: 1930, *The Right and the Good*. Oxford: Oxford University Press.
[10] Sinnott-Armstrong, W.: 1988, *Moral Dilemmas*. Oxford: Basil Blackwell.
[11] Snare, F.: 1974, 'The Definition of Prima Facie Duties'. *Philosophical Quarterly*, 24, 235-44.
[12] Vallentyne, P.: 1992, 'Moral Dilemmas and Comparative Conceptions of Morality'. *Southern Journal of Philosophy*, 30, 117-24.
[13] Williams, B.: 1973, 'A Critique of Utilitarianism'. In J. J. C. Smart and Bernard Williams (eds.) *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 77-150.
[14] Zimmerman, M. J.: 1986, 'Subsidiary Obligation'. *Philosophical Studies*, 50, 65-75.